Simpler Transfer Learning (Using "Bellwethers")

Rahul Krishna, *Member, IEEE,* Tim Menzies, *Member, IEEE*

Abstract—Transfer learning has been the subject of much recent research. In practice, that research means that the models are unstable since they are continually revised whenever new data arrives. This paper offers a very simple "bellwether" transfer learner. Given N datasets, we find which one produces the best predictions on all the others. This "bellwether" dataset is then used for all subsequent predictions (when its predictions start failing, one may seek another bellwether). Bellwethers are interesting since they are very simple to find (wrap a for-loop around standard data miners). They simplify the task of making general policies in software engineering since as long as one bellwether remains useful, stable conclusions for N datasets can be achieved by reasoning over that bellwether. This paper shows that this bellwether approach works for multiple datasets from various domains in SE. From this, we conclude that (1) bellwether method is a useful (and simple) transfer learner; (2) Unlike bellwethers, other complex transfer learners *do not* generalized to all domains in SE; (3) "bellwethers" are a *baseline method* against which future transfer learners should be compared; (4) When building increasingly complex automatic methods, researchers should pause and compare more sophisticated method against simpler alternatives.

Index Terms—Defect Prediction, Bad smells, Issue Close Time, Effort Estimation, Prediction.

arXiv:1703.06218v1 [cs.SE] 17 Mar 2017

1 Introduction

Researchers and industrial practitioners routinely make extensive use of software analytics. It has be applied in a myriad different ways. For example, to estimate how long it would take to integrate the new code [1], where bugs are most likely [2], [3], or how long it will take to develop this code [4], [5]. In fact, it is now routine for practitioners to treat every planned feature as an "experiment" [6]. In this approach, key performance metrics are carefully monitored and analyzed to judge each proposed feature. Even simple design decisions such as the color of a link are chosen by experiments [7]. Large organizations like Microsoft routinely practice data-driven policy development where organizational policies are learned from an extensive analysis of large datasets collected from developers [8], [9]. For more examples of software analytics, see Figure 1 and [10], [11].

A premise with most of the prior work on software data analytics is that there exists data from which we can learn models. This premise is not always satisfied. For example, consider a new project based on a technology that, previously, has not been used at an organization. This organization would then lack the data necessary to assess say the impact of using the technology.

When local data is scarce, sometimes it is possible to use data collected from other projects either at the local site, or other sites. That is, when building software quality predictors, it might be best to look at more than just the local data. Recent research has been exploring the problem of *transferring* data from one project to another for the purposes of data analytics. Most research has focused on two methodological variants of transfer learning: (a) dimensionality transform based techniques by Nam, Jing et al. [24], [17], [25] and (b) the similarity based approaches of Kocaguneli, Peters and Turhan et al. [26], [27], [28], [29]. In both approaches, when new code modules are created, one can comment on code quality using examples taken from other similar projects.

Transfer learning can successfully build new models in new projects using data from old projects [24], [17], [25], [26], [27], [28], [29]. However, there is a major unresolved problem with the current generation of transferring learning techniques: both dimensionality transform and similarity based approaches lack any inference of conclusion stability. Indeed, as the results of this study show, conclusions are rather unstable. Conclusions such as the efficacy of a transfer learner derived from one domain (defect prediction) do not translate well to other domains (code smells, issue lifetime, and effort estimation).

The primary cause of this instability is the constant influx of new data. When this new data arrives, it necessitates updating the transfer learners to account for additional information which may/may not be useful. In a recent study, Rahman et al. [30] warn that if quality predictors are always being updated based on the specifics of new data, then those new predictors may suffer from over-fitting. Such over-fitted models are "brittle" in the sense that they can undergo constant changes when new data arrives. That is:

> *When learning from all available data , then what we learn may be always changing whenever the available data is changed.*

Such updates are very common and occur when considering newly constructed code modules or when we learn using data from other, newly available, projects (for details on this, see §2.1 and §3.4).

Conclusion instability is unsettling for software project managers struggling to find general policies. Such instability prevents project managers offering clear guidelines on many issues including (a) when a certain module should be inspected; (b) when modules should be refactored; (c) where to focus expensive testing procedures; (d) what return-on-investment might we expect due to decreased defects after purchasing an expensive tool; etc.

How to support those managers, who seek stability in their conclusions, while also allowing new projects to take full benefit from data arriving from all the other projects constantly being completed by other programmers? Perhaps if we cannot *generalize* from all data, a more achievable goal is to *stabilize* the pace of conclusion change. While it may be a fool's errand to wait for eternal and global SE conclusions, one possible approach is for organizations to declare some prior project

The authors are with the Department of Computer Science, North Carolina State University, USA. E-mail: i.m.ralk@gmail.com, tim@menzies.us
Manuscript received XXXX XX, 20XX; revised XXXX XX, 20XX.

Effort estimation	It is possible to predict with surprising accuracy the development effort of traditional projects [12] or agile software developments [13].
Integrated risk management	Czerwonka et al. show that data miners can peek at code being checked into large projects [1] to accurately predict (a) how risky is the fix that is about to make? (b) which fewest tests might find most potential defects? (c) what new tests are required to assess this code? (d) what dependent parts of the system need to be retested?
Comparable to static analysis	Rahman et al. report that in terms of finding detects, the models learned from data miners work as well as far more elaborate and expensive static code analysis tools [14].
Security vulnerabilities	Shin and Williams report that data mining methods developed for defect prediction can also be used to find software security vulnerabilities [15].
Transfer learning	Automatic tools can transfer lessons between different projects in different countries [16], even when those projects collect different kinds of data [17].
Open source to proprietary projects	Lessons learned about software quality can even be automatically transferred between open source and proprietary projects [18].
Defect localization	Ostrand and Weyuker [2] report that data miners can find 20% of the code contain 80% of a program's defects. Other studies report similar conclusions (e.g. Turhan et al [19]'s data miners could localize 86% of the defects in 25% of the code).
Resource management	Tosen et al. report that if project teams focus on those areas found by data miners, they can half their QA costs [20];
Social patterns	Data miners can find programmer interaction patterns leading to slow fix times [21].
Project management	Data miners can also identify patterns in programmer's development processes [22], [23] that lead to less or more software errors.

Fig. 1: Some sample successes of software analytics. All these results require access to data.

as the *"bellwether"* [1] that offers predictions that generalize across *N* projects.

In this paper, we address this issue by defining and distinguishing the *bellwether effect* from the *bellwether method*:

- *The bellwether effect* states that when a community of programmers work on a set of projects, then within that community there exists one exemplary project, called the bellwether, which can define quality predictors for the other projects.
- *The bellwether method* searches for that exemplar project to construct a *simple* quality predictor. This quality predictor is then applied to all future data generated by that community.

The rest of this paper explores bellwethers from the context of: (a) Code smells; (b) Software effort estimation; (c) Estimation of issue close time; and (d) Defect prediction. After some background notes, as well as an explanation of the bellwether method, we ask and answer the following four research questions:

- **RQ1: How prevalent are "Bellwether" datasets?**
  To answer this question, we use the definition of *the bellwether effect* from above. With this we explore four sub-domains within software engineering namely, defect prediction, effort estimation, issue lifetime estimation, and detection of code smells. Each sub-domain contains multiple "communities" of datasets. In a result consistent with bellwethers being *very* prevalent, we find that all these domains have a bellwether; i.e. a single dataset from which a superior quality predictor can be generated from the rest of that community.
- **RQ2: How well do transfer learners perform across different domains?**
  To answer this question, we compare a so-called "naive" learner generated from the bellwether dataset to the predictions generated from other transfer learning methods. This naive learner performs no transformation on the data (unlike other transfer learners). It uses the data, as is. We search for an exemplar project, i.e. *Bellwether dataset*, and use it to construct quality estimation models for all future data generated by that community. We refer to the use of a naive learner+bellwether dataset as the *Bellwether Method*. When new code modules are created, this and other transfer learning approaches comment on code quality using examples taken from similar projects. Our simple *bellwether method's* predictions were observed to be superior than those of other transfer learners in two domains: effort estimation and code smell detection. *Bellwether method's* predictions were a close second in defect prediction and issue lifetime.
- **RQ3: How does the bellwether dataset fare against local models?**
  The alternate to transfer learning is to just use the local data to build a quality predictor. To answer this research question, we compare the predictions from the bellwether to predictions from just the local data. In our experiments, the bellwether predictions proved to be *better* than those generated from the local data.
- **RQ4: How much data is required to find the bellwether dataset?**
  Our proposal to find bellwethers is to compare the performance of pairs of datasets from different projects in a round robin fashion. A natural question that arises from this experimental approach is RQ4. Our experiments show that program managers need not wait very long to find their bellwethers – a few dozen samples usually are sufficient for creating and testing candidate bellwethers.

1. According to the Oxford English Dictionary, the "bellwether" is the leading sheep of a flock, with a bell on its neck.

1.1 Contributions

This work presents a significant extension to our initial findings on bellwethers [31]. A brief summary of the contributions of this paper are listed below:

1) **Rigorous Comparison of Transfer Learners:** Most

reported results in transfer learning results presented so far are inconsistent in that they use different datasets and different oracles to measure their performance. Further, their results offer no baseline comparisons. This study compares four state-of-the-art transfer learners:
   - The Bellwether Method [31];
   - Transfer Component Analysis (referred to henceforth as TCA+) [24];
   - Transfer Naive Bayes (hereafter referred to as TNB) [32]; and
   - Value Cognitive Boosting Learner [33].
2) **Generalizing Transfer Learners:** We explore the existence of Bellwethers in domains other than those traditionally used by researchers conducting transfer learning. Based on our reading of the literature, transfer learning has been shown to work relatively well for defect prediction. In this work we attempt to extend these algorithms to other domains to validate their usefulness. In specific, we study the effectiveness of Transfer Learning in:
   - Code smells detection (specifically God Class and Feature Envy);
   - Effort estimation;
   - Issue lifetime estimation; and
   - Defect Prediction.
3) **Replication Package** We note that for readers this work who wish to replicate our findings, we have made available a replication package at https://goo.gl/jCQ1Le. The replication package consists of all the datasets used in this paper, in addition to mechanisms for computation of other statistical measures.

1.2 Connection to Prior work

This paper extends the prior publication as follows:

1) This paper shows that bellwethers are an effective transfer learning method for
   - code smells;
   - effort estimation;
   - issue lifetime estimation;
   - defect prediction.

   Note that our prior publication only explored bellwethers for only one of these domains (i.e. defect prediction).
2) To the best of our knowledge, this is the first report to undertake a case study of all the state-of-the-art transfer learners and validate their usability in domains other than defect prediction.
3) While our previous work introduced *bellwethers* as a potential solution to the conclusion instability problem. The findings of this paper offer strong evidence to support this claim.
4) Our previous study compared the *bellwether method* with two relevancy based transfer learners. This paper is a more extensive case study with three new transfer learners: TCA+, Transfer Naive Bayes, and Value-Cognitive Boosting SVM.

From the above, we conclude that the original motivation for transfer learning in SE might have been misguided. Initial experiments with transfer learning in SE built quality predictors from the *union* of data taken from multiple projects. That approach lead to some very poor results so researchers turned to *relevancy filters* to find what small subset of the data was relevant to the current problem [28]. These relevancy filters generated adequate predictions but introduced the instability problem that initially motivated our previous paper [31]. Our bellwether results in the previous work suggested that relevancy filtering would never have been necessary in the first place if researchers had instead hunted for bellwethers. The findings of this work further corroborate that claim. In addition, we found that all of the recent research in transfer learning have placed much emphasis on defect prediction. While the claims about transfer learning holds true in defect prediction, it does not apply to other domains in software engineering.

2 Background

This section provides a brief background on transfer learning and bellwethers.

2.1 Transfer Learning

When there is insufficient data to apply data miners to learn defect predictors, *transfer learning* can be used to transfer lessons learned from other *source* projects S to the *target* project T.

Initial experiments with transfer learning offered very pessimistic results. Zimmermann et al. [34] tried to port models between two web browsers (Internet Explorer and Firefox) and found that cross-project prediction was still not consistent: a model built on Firefox was useful for Explorer, but not vice versa, even though both of them are similar applications. Turhan's initial experimental results were also very negative: given data from 10 projects, training on $S = 9$ source projects and testing on $T = 1$ target projects resulted in alarming high false positive rates (60% or more). Subsequent research realized data had to be carefully sub-sampled and possibly transformed before quality predictors from one source to target. That work can be divided two ways:

- *Homogeneous* vs *heterogeneous*;
- *Similarity* vs *dimensionality transform*.

Homogeneous, heterogeneous transfer learning operates on source and target data that contain the *same, different* attribute names (respectively). This paper focuses on homogeneous transfer learning, for the following reason. As discussed in the introduction, we are concerned with an IT manager trying to propose general policies across their IT organization. Organizations are defined by what they do—which is to say that within one organization there is at some overlap in task, tools, personnel, and development platforms. This overlap justifies the use of lessons derived from transfer learning.

Hence, all our dataset contain overlapping attributes. In our case these attributes are the metrics gathered for each of the projects. As evidence for this, the datasets explored in this paper fall into 4 domains; each domain contains so called "communities" of data sets. Each dataset within a community share the same attributes (see Figure 3).

As to other kinds of transfer learning, *similarity* approaches transfer some subset of the rows or columns of data from source to target. For example, the Burak filter [28] builds its training sets by finding the $k = 10$ nearest code modules in S for every $t \in T$. However, the Burak filter suffered from the all too common instability problem (described in the introduction: whenever the source or target is updated, data miners will learn a new model since different code modules will satisfy the $k = 10$ nearest neighbor criteria). Other researchers [26], [27] doubted that a fixed value of k was appropriate for all data. That work recursively bi-clustered the source data, then pruned the cluster sub-trees with greatest "variance" (where the "variance" of a sub-tree is the variance of the conclusions in its leaves). This method combined row selection with row pruning (of nearby rows with large variance). Other similarity methods [35] combine domain knowledge with automatic processing: e.g. data is partitioned using engineering judgment before automatic tools cluster the data. To address variations of software metrics between different projects, the original metric values were discretized by rank transformation according to similar degree of context factors.

Similarity approaches uses data in its raw form and as highlighted above, it suffers from instability issues. This prompted research on *Dimensionality transform* methods. These methods manipulate the raw source data until it matches the target. In the case of defect prediction, a "dimension" was one of the static code attributes of Figure 4.

An initial attempt on performing transfer learning with *Dimensionality transform* was undertaken by Ma et al. [32] with an algorithm called transfer naive Bayes (TNB). This algorithm used information from all of the suitable attributes in the training data. Based on the estimated distribution of the target data, this method transferred the source information to weight instances the training data. The defect prediction model was constructed using these weighted training data.

Nam et al. [24] originally proposed an transform-based method that used TCA based dimensionality rotation, expansion, and contraction to align the source dimensions to the target. They also proposed a new approach called TCA+, which selected suitable normalization options for TCA.

It is worth noting that researches failed to address the imbalanced classes in datasets they studied. When a dataset is gathered the samples in them tend to be skewed toward clean examples. A systematic literature review on software defect prediction carried out by Hall et al. [36] indicated that data imbalance with regard to specific classification methods may be connected to poor performance. They also suggested more studies be aware of the need to deal with data imbalance. More importantly, they assert that the performance measures chosen can *hide* the impact of imbalanced data on the real performance of classifiers.

Another approach proposed by Ryu et al. [33] showed that using Boosting-SVM combined with class imbalance learner can be used to address skewed datasets. They showed improved performance compared to TNB. More recently, in our previous work [37], we showed that a very simplistic transfer learner can be developed using the "bellwether" dataset with Random Forest. We reported highly competitive performance scores.

When there are no overlapping attributes (in heterogeneous transfer learning) Nam et al. [17] found they could dispense with the optimizer in TCA+ by combining feature selection on the source/target following by a Kolmogorov-Smirnov test to find associated subsets of columns. Other researchers take a similar approach, they prefer instead a canonical-correlation analysis (CCA) to find the relationships between variables in the source and target data [25].

Our reading of the literature suggests a surprising lack of consistency in the choice of datasets, learning methods, and statistical measures while reporting results of transfer learning. Further, there was no baseline approach to compare the algorithms against. This partly motivated our study.

2.2 Bellwethers in Software Engineering

Section §2.1 sampled some of the work on transfer learning in software engineering. This rest of this paper asks the question "is the complexity of §2.1 really necessary?"

To answer this question, we had previously proposed a process called the "Bellwether Method" [31]. We now have updated this keeping in view our findings. We propose a framework that assumes some software manager has a watching brief over N projects (which we will call the *community* "C"). As part of those duties, they can access issue reports and static code attributes of the community. Using that data, this manager will apply the a framework described in Figure 2 which comprises of three operators– DISCOVER, TRANSFER, MONITOR.

Note the simplicity of this approach– just wrap a for-loop around some data miners. Note also that these steps usually use *none* of the machinery described in §2.1.

An additional benefit of this DISCOVER-TRANSFER-MONITOR methodology is the ability to optionally replace the naive learner in the TRANSFER stage with any other transfer learner. For instance we experiment with three other transfer learners: TCA+, TNB, VCB. It is worth mentioning, the Naive transfer learner, was first reported in our prior work [37] as the "bellwether method".

A surprising outcome of such a study was the discovery of "Bellwether Effect". In the context of software engineering datasets it was found that there existed an exemplary dataset which could be used to train quality estimation models. The models trained on these exemplar datasets were capable of making very accurate quality estimations. The prevalence of the so called bellwether effect was not known in domains other than defect prediction. Hence this study.

3 Target Domains

This section describes the four domains in which we test our bellwether method.

3.1 Code Smells

According to Fowler [38], bad smells (a.k.a. code smells) are "a surface indication that usually corresponds to a deeper problem". Studies suggest a relationship between code smells and

Figure 2.A: Discover
Using historical data, check if the community has bellwether. See if data miners can predict for smells, issue lifetimes, or smells, given the code attributes.

- For all pairs of data from projects $P_i, P_j \in C$;
- Construct a quality predictor with data from project P_i and predict for issues in P_j using a quality predictor learned.
- Report a bellwether if one P_i generates the most accurate predictions in a majority of $P_j \in C$.

```
def discover(datasets):
"""
Identify Bellwether Datasets
"""
  for data_1, data_2 in datasets:
    def train(data_1):
     "Construct quality predictor"
     return predictor

    def predict(data_1):
      "Predict for quality"
      return predictions

    def score(data_1, data_2):
      "Return accuracy of Prediction"
      return accuracy(train(data_1), test(data_2))

  "Return data with best prediction score"
```

Figure 2.B: Transfer *Using the bellwether, construct a transfer learner.* That is,

- having discovered the bellwether ...
- construct a transfer learner on the bellwether data
- now apply it to future projects.

```
def transfer(datasets):
 """
 Transfer Learning with Bellwether Dataset
 """

 bellwether = discover(datasets)

 def learner(data):
   """Construct Transfer Learner, using:
     1. TCA+; 2. TNB; 3. VCB; 4. Naive
   """

 def apply_learner(datasets, learner):
   """
   Apply transfer learner
   """
   model = learner(bellwether)
   for data in datasets:
     if data != bellwether:
         train(model)
         test(data)
         yield score(model, data)
```

Figure 2.C: Monitor *Keep track of the performance of Bellwethers for transfer learning.* That is,

- If the transfer learner constructed in TRANSFER starts to fail, ...
- Go back to DISCOVER and update the bellwether.

```
def transfer(datasets):
     """
     Transfer Learning with Bellwether Dataset
     """

     def fails(data):
         "Return True if perdictions deteriorate"
```

Fig. 2: The Bellwether Framework

poor maintainability or defect proneness [39], [40], [41] and therefore, smell detection has become an established method to discover source code (or design) problems to be removed through refactoring steps, with the aim to improve software quality and maintenance. Consequently, code smells are captured by popular static analysis tools, like PMD[2], CheckStyle[3], FindBugs[4], and SonarQube[5]. However, the agreement among these detectors is scarce and this greatly impacts the results of code smell detection. The results provided by different detectors is usually different [42]. An extensive benchmark for comparison of these results is not yet available.

Most detection tools for code smells make use of so called "detection rules". Usually, these detection rules are based only on the computation of a set of metrics, e.g., well-known object-oriented metrics. These metrics are then used to set some thresholds for the detection of a code smell. Often, these thresholds are set with some "one size fits all" process. By even slightly changing these threshold values, the number of detected smells can increase or decrease accordingly. This may lead to far too many false positives making it difficult for practitioners to refactor code [37].

Recently, researchers have eschewed rule based methodologies in favor machine learning tools to classify code smells from examples, easing the build of automatic code smell detectors, thereby providing a better-targeted detection. Kreimer [43] proposes an adaptive detection to combine known methods for finding design flaws Large Class and Long Method on the basis of metrics with learning decision. Khomh et al. [44] propose a Bayesian approach to detect occurrences of the Blob antipattern on open-source programs (GanttProject v1.10.2 and Xerces v2.7.0). Khomh et al. [45] also present BDTEX (Bayesian Detection Expert), a GQM approach to build Bayesian Belief Networks from the definitions of antipatterns and validate BDTEX with Blob, Functional Decomposition, and Spaghetti Code antipatterns on two open-source programs. Yang et al. [46] study the judgment of individual users by applying machine learning algorithms on code clones. These studies were not included in our comparison as the data was not readily available for us to reuse.

More recently, Fontana et al. [47] in their study of several code smells, considered 74 systems for their analysis and their validation and they experimented with 16 different machine learning algorithms. They made available their dataset, which we have adapted for our applications in this study.

As a result of the ongoing research on the application of machine learning for code smell detection problems, a general formalization of the input and output of the learning algorithm and a selection of data and algorithms to be used in the experimentation has become widely accepted. This formalization is strikingly similar to that of software defect prediction, thus opening the possibility of using transfer learning to transfer useful knowledge between projects in detecting and addressing code smells.

2. https://github.com/pmd/pmd
3. http://checkstyle.sourceforge.net/
4. http://findbugs.sourceforge.net/
5. http://www.sonarqube.org/

Defect

Community	Dataset	# of instances		# metrics	Nature
		Total	Bugs (%)		
AEEEM	EQ	325	129 (39.81)	61	Class
	JDT	997	206 (20.66)		
	LC	399	64 (9.26)		
	ML	1826	245 (13.16)		
	PDE	1492	209 (13.96)		
Relink	Apache	194	98 (50.52)	26	File
	Safe	56	22 (39.29)		
	ZXing	399	118 (29.57)		
Apache	Ant	1692	350 (20.69)	20	Class
	Ivy	704	119 (16.90)		
	Camel	2784	562 (20.19)		
	Poi	1378	707 (51.31)		
	Jedit	1749	303 (17.32)		
	Log4j	449	260 (57.91)		
	Lucene	782	438 (56.01)		
	Velocity	639	367 (57.43)		
	Xalan	3320	1806 (54.40)		
	Xerces	1643	654 (39.81)		

Code Smells

Community	Dataset	# of instances		# metrics	Nature
		Samples	Smelly (%)		
Feature Envy	wct	25	18 (72.0)	83	Method
	itext	15	7 (47.0)		
	hsqldb	12	8 (67.0)		
	nekohtml	10	3 (30.0)		
	galleon	10	3 (30.0)		
	sunflow	9	1 (11.0)		
	emma	9	3 (33.0)		
	mvnforum	9	6 (67.0)		
	jasml	8	4 (50.0)		
	xmojo	8	2 (25.0)		
	jhotdraw	8	2 (25.0)		
God Class	fitjava	27	2 (7.0)	62	Class
	wct	24	15 (63.0)		
	xerces	17	11 (65.0)		
	hsqldb	15	13 (87.0)		
	galleon	14	6 (43.0)		
	xalan	12	6 (50.0)		
	itext	12	6 (50.0)		
	drjava	9	4 (44.0)		
	mvnforum	9	2 (22.0)		
	jpf	8	2 (25.0)		
	freecol	8	7 (88.0)		

Effort Estimation

Community	Dataset	Samples	Range (min-max)	# metrics
Effort	coc10	95	3.5 - 2673	24
	nasa93	93	8.4 - 8211	
	coc81	63	5.9 - 11400	
	nasa10	17	320 - 3291.8	
	cocomo	12	1 - 22	

Issue Lifetime

Community	Dataset	# of instances		# metrics
		Total	Closed (%)	
camel	1 day	5056	698 (14.0)	18
	7 days		437 (9.0)	
	14 days		148 (3.0)	
	30 days		167 (3.0)	
cloudstack	1 day	1551	658 (42.0)	18
	7 days		457 (29.0)	
	14 days		101 (7.0)	
	30 days		107 (7.0)	
cocoon	1 day	2045	125 (6.0)	18
	7 days		92 (4.0)	
	14 days		32 (2.0)	
	30 days		45 (2.0)	
node	1 day	2045	125 (6.0)	18
	7 days		92 (4.0)	
	14 days		32 (2.0)	
	30 days		45 (2.0)	
deeplearning	1 day	1434	931 (65.0)	18
	7 days		214 (15.0)	
	14 days		76 (5.0)	
	30 days		72 (5.0)	
hadoop	1 day	12191	40 (0.0)	18
	7 days		65 (1.0)	
	14 days		107 (1.0)	
	30 days		396 (3.0)	
hive	1 day	5648	18 (0.0)	18
	7 days		22 (0.0)	
	14 days		58 (1.0)	
	30 days		178 (3.0)	
ofbiz	1 day	6177	1515 (25.0)	18
	7 days		1169 (19.0)	
	14 days		467 (8.0)	
	30 days		477 (8.0)	
qpid	1 day	5475	203 (4.0)	18
	7 days		188 (3.0)	
	14 days		84 (2.0)	
	30 days		178 (3.0)	

Fig. 3: Datasets from 4 chosen domains.

3.1.1 Code Smells Data

The datasets were generated by Fontana et al. [47] using the Qualitas Corpus (QC) of systems [48]. The Qualitas corpus is composed of 111 systems written in Java, characterized by different sizes and belonging to different application domains. Fontana et al. selected a subset of 74 systems for their analysis. The authors computed a large set of object-oriented metrics. The selected metrics belonged to class, method, package and project level. Additionally, standard metrics covering different aspects of the code, i.e., complexity, cohesion, size, coupling. A detailed list of metrics and their definitions are available in appendices of [47].

The code smells repository we use comprises of 22 datasets for two different code smells: Feature envy and God Class. The God Class code smell refers to classes that tend to centralize the intelligence of the system. This tends to be: (1) complex; (2) have too much code; (3) use large amounts of data from other classes; and (4) implement several different functionalities. Feature Envy tends to: (1) use many attributes of other classes (considering also attributes accessed through accessor methods); (2) use more attributes from other classes than from its own class; and (3) use many attributes from few different classes. The God Class is a class level smell and Feature Envy is a method level design smell.

The number of samples in these datasets are particularly small. For our analysis, we retained only datasets with at least

Size		Complexity		Cohesion		Coupling		Encapsulation		Inheritance	
Label	Description	Label	Description	Label	Description	Label	Description	Label	Description	Label	Description
LOC	Lines Of Code	CYCLO	Cyclomatic Complexity	LCOM	Lack of Cohesion	FANOUT/IN	Fan Out/In	LAA	Locality of Attribute Accesses	DIT	Depth of Inheritance Tree
LOCNAMM	LOC (without accessor or mutator)	WMC	Weighted Methods Count	TCC	Tight Class Cohesion	ATFD	Access to Foreign Data	NOAM	Number of Accessor Methods	NOI	Number of Interfaces
NOM	No. of Methods	WMCNAMM	Weighted Methods Count (without accessor or mutator)	CAM	Cohesion Among classes	FDP	Foreign Data Providers	NOPA	Number of Public Attribute	NOC	Number of Children
NOPK	No. of Packages	AMW	Average MethodsWeight			RFC	Response for a Class			NMO	Number of Methods Overridden
NOCS	No. of Classes	AMWNAMM	Average Methods Weight (without accessor or mutator)			CBO	Coupling Between Objects			NIM	Number of Inherited Methods
NOMNAMM	Number of Not Accessor or Mutator Methods	MAXNESTING	Max Nesting			CFNAMM	Called Foreign Not Accessor or Mutator Methods			NOII	Number of Implemented Interfaces
NOA	Number of Attributes	CLNAMM	Called Local Not Accessor or Mutator Methods			CINT	Coupling Intensity				
		NOP	Number of Parameters			MaMCL	Maximum Message Chain Length				
		NOAV	Number of Accessed Variables			MeMCL	Mean Message Chain Length				
		ATLD	Access to Local Data			CA/CE/IC	Afferent/ Efferent/ Inheritance coupling				
		NOLV	Number of Local Variable			CM	Changing Methods				
		WOC	Weight Of Class			CBM	Coupling between Methods				
		MAX_CC/AVG_CC	Maximum/ Average McCabe								

Fig. 4: Static code metrics used in defects and code smells data sets.

Commit	Comment	Issue
nCommitsByActorsT	meanCommentSizeT	issueCleanedBodyLen
nCommitsByCreator	nComments	nIssuesByCreator
nCommitsByUniqueActorsT		nIssuesByCreatorClosed
nCommitsInProject		nIssuesCreatedInProject
nCommitsProjectT		nIssuesCreatedInProjectClosed
		nIssuesCreatedProjectClosedT
		nIssuesCreatedProjectT
Misc.	nActors, nLabels, nSubscribedByT	

Fig. 5: Metrics used in issue lifetimes data.

Personnel		Product		System		Other	
Label	Description	Label	Description	Label	Description	Label	Description
ACAP	Analyst Capability	CPLX	Prod. Complexity	DATA	Database size	DOCU	Documentation
APEX	Applications Exp.	SCED	Dedicated Schedule	PVOL	Platform volatility	TOOL	Use of software tools
LEXP	Language Exp.	SITE	Multi-side dev.	RELY	Required Reliability		
MODP	Modern Prog. Practices	TURN	turnaround time	RUSE	Required Reuse		
PCAP	Programmer Capability			STOR	% RAM		
PLEX	Platform Exp.			TIME	% CPU time		
VEXP	Virtual Machine Exp.			VIRT	Machie volatility		
PCON	Personnel Continuity						

Fig. 6: Metrics used in effort estimation dataset.

8 samples. All our transfer learners required that at least 8 samples be present in order for them to function reliably. This lead us to a total of 22 datasets shown in Figure 3.

3.2 Issue Lifetime Estimation

Open source projects use issue tracking systems to enable effective development and maintenance of their software systems. Typically, issue tracking systems collect information about system failures, feature requests, and system improvements. Based on this information and actual project planing, developers select the issues to be fixed. Predicting the time it may take to close an issue has multiple benefits for the developers, managers, and stakeholders involved in a software project. Predicting issue lifetime helps software developers better prioritize work; helps managers effectively allocate resources and improve consistency of release cycles; and helps project stakeholders understand changes in project timelines and budgets. It is also useful to be able to predict issue lifetime specifically when the issue is created. An immediate prediction can be used, for example, to auto-categorize the issue or send a notification if it is predicted to be an easy fix.

As an initial attempt, Panjer [49] used logistic regression models to classify bugs as closing in 1.4, 3.4, 7.5, 19.5, 52.5, and 156 days, and greater than 156 days. He was able to achieve an accuracy of 34.9%. Giger et al. [50] used models constructed with decision trees to predict for issue lifetimes in Eclipse, Gnome, and Mozilla. They were able obtain a peak precision of 65%. By dividing time in 1, 3, 7, 14, 30 days and using static features they reached a peak precision and recall of 63% and 65% respectively.

In their ICSE '13 paper, Zhang et al. [51] developed a comprehensive, but complex, system to predict lifetime of issues. They used a Markov model with a kNN-based classifier to perform their prediction. More recently, Rees-Jones et al [52] show that using Hall's CFS feature selector and C4.5 decision tree learner a very reliable prediction of issue lifetime could be made circumventing much of the complex mathematical formalisms.

In our work, we approach issue lifetime estimation from the perspective of transfer learning. As previously highlighted, much success has been reported with transfer learners in defect prediction, and we investigate if features that predict for issue close times within different projects can be effectively transferred between each other.

3.2.1 Issue Lifetime Data

Figure 3 shows a list of 8 projects used to study issue lifetimes. These projects were selected by our industrial partners since they use, or extend, software from these projects. It forms a part of an ongoing study on prediction of issue lifetime by Rees-Jones et al. [52]. The raw data dumps were in the form of JSON files which came in the form of commit data, issues, and code contributors from GitHub and JIRA projects, the authors extracted the issue datasets from them with a minimum of 1,434 issues, maximum of 12,191, and median of 5,267 issues per dataset.

The authors note that one issue in preparing their data was a small number of *sticky* issues. They define sticky issues as one which was not yet closed at the time of data collection. When faced with a similar situation kikas et al. [53] handled it by approximating the lifetime to be a chosen set date in the future. Although this was innovative, such an approach is somewhat subjective. Hence, as recommended by Rees-Jones et al. [52], we removed these sticky issues from our datasets. It should be noted that, when studying local learning for RQ2, (where training and test data come from the same source), as noted by Rees-Jones et al. it is unknown whether or not this handling of "sticky" issues is to preferred to that of Kikas et al. However, for the sake of consistency, we follow the recommendations of [52] and remove these issues from the datasets.

In raw form, the data consisted of sets of JSON files for each repository, each file contained one type of data regarding the software repository (issues, commits, code contributors, changes to specific files), with one JSON file joining all of the data associated with a commit together: the issue associated with a commit, commit time, the magnitude of the commit, and other information. In order to extract data specific to issue lifetime, we did similar preprocessing and feature extraction on the raw datasets as suggested by [52].

3.3 Effort Estimation

Kitchenham et al. [54] reviewed 7 published transfer studies in effort estimation. They found that in most (10) cases, transferred data generated worse predictors than using within-project information. Similarly, Ye et al. [55] report that the tunings to Boehms COCOMO model have changed radically for new data collected in the period 2000 to 2009.

The key challenge with effort estimation is the nature of the dataset. Firstly, while defect prediction datasets often store several thousand samples of defective and non-defective samples, effort data is usually smaller with only a few dozen samples at most. Secondly, unlike defect dataset or code smells, effort is measured using, say *man-hours*, which is a continuous variable. These differences requires us to significantly modify existing transfer learning techniques to accommodate this kind of data.

It is worth noting that, Kocaguneli et al. [27] used analogy-based effort estimation with relevancy filtering using a method called TEAK for studying transfer learning in effort estimation. He found that it out-performs other approaches such as linear regression, neural networks, and traditional analogy-based reasoners. Since then, however, newer more sophisticated transfer learners have been introduced. Moreover, Krishna et al. [31] suggest that relevancy filtering (for defect prediction tasks) would never have been necessary in the first place if researchers had instead hunted for bellwethers. Therefore, in this paper, we revisit transfer learning in effort estimation keeping in mind the changing trends.

3.3.1 Effort Estimation dataset

We consider effort estimation data expressed in terms of the COCOMO ontology: 23 attributes describing a software project, as well as aspects of its personnel, platform, and system features (see Figure 6 for details). The data is gathered using Boehm's 2000 COCOMO model. The data was made

available by Menzies et al. [56] who show that this model works better than (or just as well as) other models they've previously studied.

In this study, we use 5 datasets shown in Figure 3. Here, COC81 is the original data from 1981 COCOMO book [57]. This comes from projects dated from 1970 to 1980. NASA93 is NASA data collected in the early 1990s about software that supported the planning activities for the International Space Station. The other datasets are NASA10 and COC05 (the latter is proprietary and cannot be released to the research community). The non-proprietary data (COC81 and NASA93 and NASA10) are available at http://tiny.cc/07wvjy.

3.4 Defect Prediction

Hall et al. [36] offers an extensive review on the defect prediction literature. For an extensive experimental comparison of different learning algorithms for defect prediction, see Lessmann et al. [58]. The rest of this section provides a brief introductory note on defect prediction.

Human programmers are clever, but flawed. Coding adds functionality, but also defects, so software will crash (perhaps at the most awkward or dangerous time) or deliver wrong functionality.

Since programming introduces defects into programs, it is important to test them before they are used. Testing is expensive. According to Lowry et al. software assessment budgets are finite while assessment effectiveness increases exponentially with assessment effort [59]. Exponential costs quickly exhaust finite resources, so standard practice is to apply the best available methods only on code sections that seem most critical. Any method that focuses on parts of the code can miss defects in other areas so some sampling policy should be used to explore the rest of the system. This sampling policy will always be incomplete, but it is the only option when resources prevent a complete assessment of everything.

One such lightweight sampling policy is defect predictors learned from static code attributes. Given software described in the attributes of Figures 4, 5, and 6, data miners can learn where the probability of software defects is highest.

The rest of this section argues that such defect predictors are *easy to use*, *widely-used*, and *useful* to use.

Easy to use: Static code attributes can be automatically collected, even for very large systems [60]. Other methods, like manual code reviews, are far slower and far more labor-intensive. For example, depending on the review methods, 8 to 20 LOC/minute can be inspected and this effort repeats for all members of the review team, which can be as large as four or six people [61].

Widely used: Researchers and industrial practitioners use static attributes to guide software quality predictions. Defect prediction models have been reported at Google [62]. Verification and validation (V&V) textbooks [63] advise using static code complexity attributes to decide which modules are worth manual inspections.

Useful: Defect predictors often find the location of 70% (or more) of the defects in code [64]. Defect predictors have some level of generality: predictors learned at NASA [64] have also been found useful elsewhere (e.g. in Turkey [19], [20]). The success of this method in predictors in finding bugs is markedly higher than other currently-used industrial methods such as manual code reviews. For example, a panel at *IEEE Metrics 2002* [65] concluded that manual software reviews can find ≈60% of defects. In another work, Raffo documents the typical defect detection capability of industrial review methods: around 50% for full Fagan inspections [66] to 21% for less-structured inspections.

Not only do static code defect predictors perform well compared to manual methods, they also are competitive with certain automatic methods. A recent study at ICSE'14, Rahman et al. [14] compared (a) static code analysis tools FindBugs, Jlint, and Pmd and (b) static code defect predictors (which they called "statistical defect prediction") built using logistic regression. They found no significant differences in the cost-effectiveness of these approaches. Given this equivalence, it is significant to note that static code defect prediction can be quickly adapted to new languages by building lightweight parsers that find information like Figure 4. The same is not true for static code analyzers– these need extensive modification before they can be used on new languages.

3.4.1 Defect datasets

The defect dataset comprises of 59 datasets grouped into 4 communities taken from previous transfer learning studies. The projects measure defects at various levels of granularity ranging from function-level to file-level. Figure 3 summarizes all the communities of datasets used in our experiments.

For the reasons discussed in §2.1, we explore homogeneous transfer learning using the attributes shared by a community. That is, this study explores intra-community transfer learning and not cross-community heterogeneous transfer learning.

The first dataset, AEEEM, was used by [17]. This dataset was gathered by D'Amborse et al. [67], it contains 61 metrics: 17 object-oriented metrics, 5 previous-defect metrics, 5 entropy metrics measuring code change, and 17 churn-of-source code metrics.

The RELINK community data was obtained from work by Wu et al. [68] who used the Understand tool [6], to measure 26 metrics that calculate code complexity in order to improve the quality of defect prediction. This data is particularly interesting because the defect information in it has been manually verified and corrected. It has been widely used in defect prediction [17][68][69][70][71].

In addition to this, we explored two other communities of datasets from the SEACRAFT repository[7]. The group of data contains defect measures from several Apache projects. It was gathered by Jureczko et al. [72]. This dataset contains records the number of known defects for each class using a post-release bug tracking system. The classes are described in terms of 20 OO metrics, including CK metrics and McCabes complexity metrics. Each dataset in the Apache community has several versions. There are a total of 38 different datasets. For more information on this dataset see [37].

6. http://www.scitools.com/products/
7. https://zenodo.org/communities/seacraft/

4 Research Questions

RQ1: How prevalent are "Bellwethers"?

If bellwethers occur infrequently, we cannot rely on them. Hence, this question explores how common are bellwethers. To this end, we applied the DISCOVER method described above to all the available datasets in each sub-domain shown in Figure 3.

For each sub-domain, we ensured that the datasets were as diverse as possible. To this end, data was gathered according to the following rules:

- The data has been used in a prior paper. Each of our datasets for defects, code smells, effort estimation, and issue lifetime estimation has been used previously, see §2 for further details;
- The communities are quite diverse; e.g. the NASA projects from the effort estimation datasets are proprietary while the others are open source projects. Similarly, The God Class is a class level smell and Feature Envy is a method level design smell.
- In addition, where relevant, the projects also vary in their granularity of data description (file, class, or function level).

RQ2: How well do transfer learners perform across different domains?

Our reading of the literature is that for homogeneous transfer learning, the current state of the art according to the literature is to use TCA+.

However, as previously noted that this result has only been validated for defect prediction and only for a limited number of datasets. In our previous work we reported that Bellwether was better than relevancy based filtering methods. Here we ask if this is true given newer transfer learning methods and different datasets.

To answer this question, we compare the "bellwether" method [31] against 3 other standard transfer learners: (1) TCA+ [24]; (2) Transfer Naive Bayes [32]; and (3) Value Cognitive Boosting [33]. In addition we modify these learners appropriately for different sub-domains under study.

RQ3: How does the bellwether fare against local models?

One premise of transfer learning is that using data from other projects is as useful, or better, than using data from the local project. This research questions tests that this premise holds for bellwethers.

To answer this question, we reflect on datasets with temporal local data. From Figure 3 we see that data from Code Smells and Effort Estimation have far to little samples to study the effect of local data. However, one of our communities in defect prediction (APACHE) comes in multiple versions. We use this to answer this research question.

We implemented TRANSFER as follows. The XALAN system has versions 2.4, 2.5, 2.6, 2.7. Each versions are historical releases where version i was written before version j where $j > i$. RQ3 was explored in this community as follows:

- The last version of each project was set aside as a hold-out.
- DISCOVER was then applied across the older versions within the community to find the bellwether.
- A defect predictor was then learned from the older data seen in the bellwether.
- The predictor was then applied to the latest data.

We compare the above to *local learning*; i.e. for each project:

- The last version of that project was set aside as a hold-out;
- The older versions of that project were then used to train a defect predictor.
- The predictor was then applied to the latest data.

Note that:

- The local learner only ever uses data from earlier in the *same project*;
- While the bellwether uses data from *any member* of the community.

RQ4: How much data is required to find the bellwether dataset?

A core process in all the above is the DISCOVER step. If this requires too much data to find bellwethers, then we may have to wait far too long to accumulate required amount of data. That would discourage developers from looking for bellwethers. They may instead prefer standard transfer learning methodologies. Hence, it is important to ask how much data is required before a community can find adequate bellwethers.

5 Methodology

5.1 Learning Methods

There are many ways to predict defects. A comprehensive study on the same was conducted by Lessmann et al. [58]. They endorsed the use of Random Forests [73] for defect prediction over several other methods. This was also true in detecting code smells [47]. When a specific transfer learner did not endorse the use of any classification/regression scheme, we used Random Forests.

Random Forests is an ensemble learning method that builds several decision trees on randomly chosen subsets of data. The final reported prediction is the mode of predictions by the trees.

It is known that the fraction of interesting samples in the training data affects the performance of predictors. Figure 3 shows that in most datasets, the percentage of interesting samples (i.e., samples that are defective, smelly, closed, or ones that require more effort) vary between 10% to 40% (except in a few, projects like log4j for instance where it is 58%). Handling this class imbalance has been shown to improve the quality of prediction.

Pelayo and Dick [74] report that the defect prediction is improved by SMOTE [75]. SMOTE works by under-sampling majority-class examples and over-sampling minority class examples to balance the training data prior to applying prediction models.

After an extensive experimentation, in this study, we randomly sub-sampled examples until the training data had only *100,50* uninteresting, interesting examples (respectively).

Important methodological notes:

1) sub-sampling was only applied to *training* data (so the test data remains unchanged).
2) Authors of several transfer learners studied here recommend using different predictors. When replicating their studies, we adhere to their recommendations.
3) SMOTE works only for classification problems (defect prediction, code smell detection, and issue lifetime prediction). When performing regression for estimation of effort, we don't apply SMOTE.

5.2 Evaluation Strategy

5.3 Evalutaion for Continuous Classes

For the effort estimation data in Figure 3, the dependent attribute is development effort, measured in terms of calendar months (at 152 hours per month, including development and management hours). For this, we use the same learning methods as in §5.1 used as a regressor instead of a classifier.

To evaluate the quality of the learners used for regression, we make use of MMRE. MMRE is computed as below:

$$MMRE = \frac{abs(actual - predicted)}{actual} \tag{1}$$

Lower values of MMRE are considered to be *better*. Note: Some researchers have endorsed the use other metrics such as standard error to measure the quality of regressor in effort estimation. We have made available a replication package (see §1.1) with this and other metrics. Interested readers are encouraged to use these.

5.4 Evalutaion for Discrete Classes

In the context of discrete classes, we define interesting and uninteresting modules. With defects, instances with one or more defects are considered "interesting". As in defect prediction so too in code smell detection (smelly samples are interesting) and issue lifetime estimation (closed issues are interesting). Prediction models are not ideal, they therefore need to be evaluated in terms of statistical performance measures.

For classification problems we construct a confusion matrix, with this we can obtain several performance measures such as: (1) *Accuracy*: Percentage of correctly classified classes (both positive and negative); (2) *Recall or pd*: percentage of the target classes (defective instances) predicted. The higher the pd, the fewer the false negative results; (3) *False alarm or pf*: percentage of non-defective instances wrongly identified as defective. Unlike pf, lower the pd better the quality; (4) *Precision*: probability of predicted defects being actually defective. Either a smaller number of correctly predicted faulty modules or a larger number of erroneously predicted defect-free modules would result in a low precision.

There are several trade-offs between the metrics described above. There is a trade-off between recall rate and false alarm rate. There is also a trade-off between precision and recall. These measures alone do not paint a complete picture of the quality of the predictor. Therefore, it is very common to apply performance metrics that incorporate a combination of these metrics. One such approach is to build a *Receiver Operating Characteristic (ROC)* curve. ROC curve is a plot of Recall versus False Alarm pairing for various predictor cut-off values ranging from 0 to 1. The best possible predictor is the one with an ROC curve that rises as steeply as possible and plateaus at pd=1.

Ideally, For each curve, we can measure the *Area Under Curve (AUC)*, to identify the best training dataset. Unfortunately, building an ROC is not straight forward in our case. We have used Random Forest for predicting defects owing to it's superior performance over several other predictors [58]. Random Forest lacks a threshold parameter, it is capable of producing just one point on the ROC curve. It is therefore not possible to compute AUC. In a previous work, Ma and Cukic [76] have shown that distance from perfect classification (ED) can be substituted for AUC in cases where a ROC curve cannot be generated. ED measures the distance between obtained (Pd, Pf) pair and the ideal point on the ROC space (1, 0), weighted by cost function θ. It is given by:

$$ED = \sqrt{\theta.(1-Pf)^2 + (1-\theta).Pd^2} \tag{2}$$

Note that for ED, the *smaller* the distance, the *better* the predictor. Alternatively, Menzies et al. [3] suggest using the "G-Score" for combining Pd and Pf. They showed that it is justifiably better than other measures when the test samples have imbalanced distribution in terms of classes. G Score is computed by measuring the harmonic mean between the Probability of True Positives (Pd) and Probability of true negatives (1-Pf). That is,

$$G = \frac{2 \times Pd \times (1-Pf)}{1 + Pd - Pf} \tag{3}$$

In our previous work, we reported only on ED-measures. However, for the sake of reliability and consistency with [3], this paper contains measures of Pd and Pf reported in terms of the G-measure.

5.5 Statistics

To overcome the inherent randomness introduced by Random Forests and SMOTE, we use 30 repeated runs, each time with a different random number seed (we use 30 since that is the minimum needed samples to satisfy the central limit theorem). The repeated runs provide us with a sufficiently large sample size to statistically compare all the datasets. Each run collects the values of Pd and Pf which are then used to estimate the G-Score using Equation 3. (Note: We refrain from performing a cross validation because the process tends to mix the samples from training data (the source) and the test data (other target projects), which defeats the purpose of this study.)

To rank these 30 numbers collected as above, we use the Scott-Knott test recommended by Mittas and Angelis [77]. Scott-Knott is a top-down clustering approach used to rank different treatments. If that clustering finds an interesting division of the data, then some statistical test is applied to

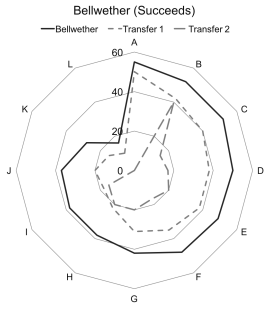

Figure 7.A: In this first example figure, the thick line shows the bellwether results and the dotted lines show results from other transfer learning methods. The key feature of this figure is that the for the data sets A,B,C,D...L, the thick Bellwether line falls outside all the rest. That is, results like this can be summarized as "bellwether performs *better* than the other transfer learning methods".

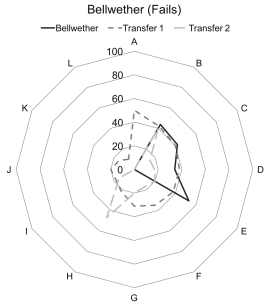

Figure 7.B: In this second example figure, for the data sets A,B,C,D...L, the thick Bellwether line often falls inside all the rest. That is, results like this can be summarized as "bellwether often performs *worse* than the other transfer learning methods".

Fig. 7: Reference charts (with contrived values) shows bellwether succeeding and failing.

the two divisions to check if they are statistically significant different. If so, Scott-Knott recurses into both halves.

To apply Scott-Knott, we sorted a list of $l = 40$ values of Equation 2 values found in $ls = 4$ different methods. Then, we split l into sub-lists m,n in order to maximize the expected value of differences in the observed performances before and after divisions. E.g. for lists l,m,n of size ls,ms,ns where $l = m \cup n$:

$$E(\Delta) = \frac{ms}{ls} abs(m.\mu - l.\mu)^2 + \frac{ns}{ls} abs(n.\mu - l.\mu)^2$$

We then apply a statistical hypothesis test H to check if m,n are significantly different (in our case, the conjunction of A12 and bootstrapping). If so, Scott-Knott recurses on the splits. In other words, we divide the data if *both* bootstrap sampling and effect size test agree that a division is statistically significant (with a confidence of 99%) and not a small effect ($A12 \geq 0.6$). For a justification of the use of non-parametric bootstrapping, see Efron & Tibshirani [78, p220-223]. For a justification of the use of effect size tests see Shepperd and MacDonell [79]; Kampenes [80]; and Kocaguenli et al. [81]. These researchers warn that even if a hypothesis test declares two populations to be "significantly" different, then that result is misleading if the "effect size" is very small. Hence, to assess the performance differences we first must rule out small effects using A12, a test recently endorsed by Arcuri and Briand [82].

6 Results

6.1 Understanding data visualization

In presenting our results we adopted a convention that includes a table accompanied by radar charts like those as shown in Figure 7. The following remarks need to be made regarding our charts and tables.

1) Each transfer learner is represented as a polygon with different line types. The best transfer learner is the polygon which encompasses the largest area.
2) datasets are arranged along the circumference. They are sorted according to their median performance score. The "best" appears on the top with the rest arranged clockwise.
3) The two figures in Figure 7 show examples where the bellwether method performs *better* or *worse* than other transfer learning method.
   - In Figure 7.A, the thick line showing the bellwether results falls *outside* all the rest. Results such as this can be summarized as bellwether method performing *better* than other transfer learners.
   - In Figure 7.B, bellwether method performs *worse* than the other transfer learners, as evidenced by the observation that the thick line showing the bellwether results falls *inside* the rest.
4) Beside every chart, the numeric results are presented in a tabular format. Here:
   - The second column labeled "source" indicates the source from which a transfer learner is built. The remaining datasets within the community are then used as target datasets.
   - The numeric values indicate the *median* performance scores (MMRE in case of effort estimation, G-score in the rest), when model is constructed with a "target" dataset and tested against all the "source" datasets, and this processes repeated 30 times for reasons discussed in §5.5.

RQ1: How prevalent is the "Bellwether Effect"?

The bellwether effect points to an exemplar dataset to construct quality estimation models form. Ideally, given a adequate transfer learner, such a dataset should produce reasonably high performance scores. Figures 8, 9,10, and 11 show the results of applying the various transfer learners to our datasets.

It is immediately noticeable that for each community there is one dataset that provides consistently better predictions when compared to other datasets. For example:

1) *Code Smells datasets:* Here Xalan is the bellwether for when predicting for the existence of God Classes; when Xalan was absent in Feature Envy, we found mvnforum to the bellwether.

	Source	**Bellwether Method**	TCA	TNB
God Class	**xalan**	**90**	75	48
	xerces	90	73	39
	hsqldb	88	0	0
	galleon	87	61	55
	wct	81	58	67
	drjava	80	58	56
	jpf	79	59	65
	mvnforum	74	43	57
	freecol	69	0	0
	fitjava	68	40	0
	itext	62	72	30
	W/T/L	**10/0/1**	1/0/10	0/0/11
Feature Envy	**mvnforum**	**92**	57	61
	galleon	84	59	0
	hsqldb	81	57	0
	jhotdraw	81	35	64
	nekohtml	81	52	57
	wct	81	47	0
	itext	74	66	0
	xmojo	74	0	0
	emma	70	74	37
	jasml	66	79	0
	sunflow	47	0	0
	W/T/L	**10/0/1**	1/0/10	0/0/11

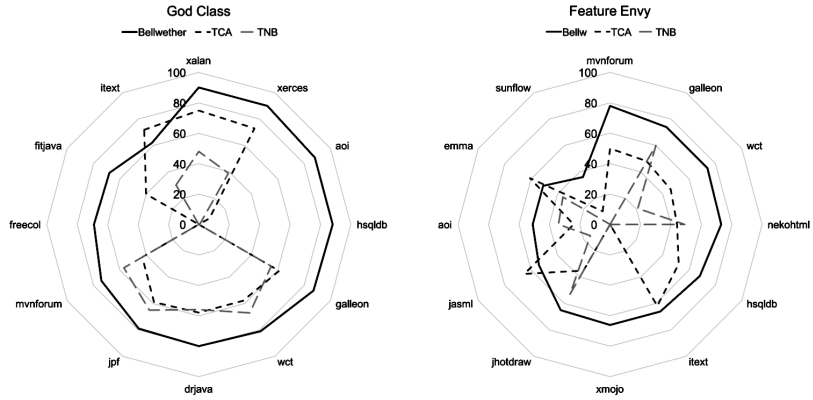

Fig. 8: Code Smells: This figure compares the prediction performance of the bellwether dataset (xalan,mvnforum) against other datasets (other rows). *Bellwether Method* against Transfer Learners (columns) for detecting code smells. The numerical value seen here are the median G-scores from Equation 2 over 30 repeats where one dataset is used as a source and others are used as targets in a round-robin fashion. Higher values are better and cells highlighted in gray produce the best Scott-Knott ranks. The last row in each community indicate Win/Tie/Loss(W/T/L). The *bellwether Method* is the overall best.

	Source	Bellwether Method	TCA+	**TNB**	VCB
1 Day	**qpid**	17	2	55	10
	node	44	24	55	17
	ofbiz	29	14	53	8
	camel	44	34	49	19
	deeplearning	51	42	42	15
	cocoon	7	6	34	13
	cloudstack	55	32	32	11
	hive	11	1	22	23
	hadoop	17	0	10	19
	W/T/L	0/0/9	0/0/9	**7/0/2**	2/0/7
7 Days	**qpid**	17	3	49	11
	camel	34	6	47	20
	cloudstack	8	27	38	7
	ofbiz	7	16	38	20
	node	15	33	36	13
	deeplearning	15	20	29	10
	cocoon	0	3	22	16
	hadoop	23	0	18	19
	hive	3	0	7	14
	W/T/L	0/0/9	0/0/9	**7/0/2**	2/0/7
14 Days	**qpid**	0	0	39	6
	cloudstack	8	8	36	8
	hadoop	0	0	31	22
	deeplearning	4	6	30	17
	camel	1	6	29	18
	cocoon	0	0	19	8
	node	5	4	16	4
	ofbiz	2	2	7	12
	hive	0	0	0	14
	W/T/L	0/0/9	0/0/9	**7/0/2**	2/0/7
30 Days	**qpid**	1	5	47	17
	cloudstack	1	13	38	19
	node	2	10	32	16
	camel	1	1	30	17
	deeplearning	1	2	29	14
	cocoon	2	1	24	12
	ofbiz	4	5	13	5
	hadoop	0	0	7	2
	hive	16	2	6	9
	W/T/L	1/0/8	0/0/9	**8/0/1**	0/0/9

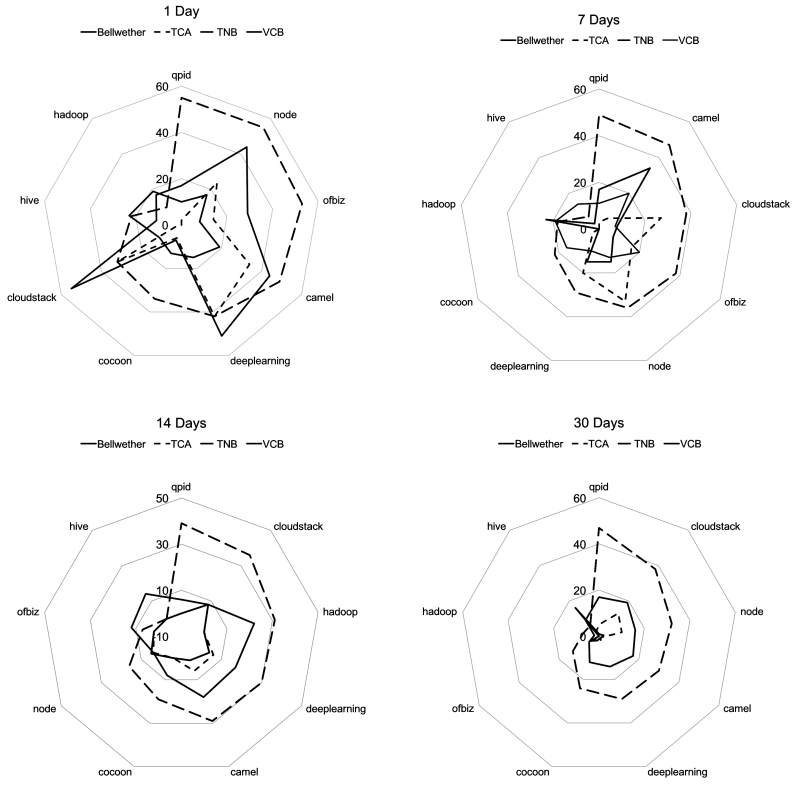

Fig. 9: Issue Lifetime: This figure compares the prediction performance of the bellwether dataset (qpid) against other datasets (rows) and various transfer learners (columns) for estimating issue lifetime. The numerical value seen here are the median G-scores from Equation 2 over 30 repeats where one dataset is used as a source and others are used as targets in a round-robin fashion. Higher values are better and cells highlighted in gray produce the best Scott-Knott ranks. The last row in each community indicate Win/Tie/Loss(W/T/L). TNB has the overall best Win/Tie/Loss ratio.

	Source	Bellwether Method	**TCA+**	TNB	VCB
Jureczko	**Lucene**	63	**69**	57	64
	Xalan	57	64	59	62
	Camel	60	63	59	44
	Velocity	58	63	51	63
	Ivy	60	62	61	48
	Log4j	60	62	58	62
	Xerces	57	54	58	65
	Ant	61	52	45	55
	Jedit	58	43	57	49
	W/T/L	2/0/7	**6/0/3**	0/0/9	1/2/06
ReLink	**Zxing**	**68**	**67**	53	64
	Safe	38	34	36	31
	Apache	31	31	32	31
	W/T/L	**0/1/1**	0/1/1	1/0/2	0/1/2
AEEEM	**LC**	**75**	**75**	73	61
	ML	73	73	67	51
	PDE	70	71	60	57
	JDT	63	64	68	53
	EQ	59	61	59	57
	W/T/L	0/2/3	**2/2/1**	1/0/4	0/0/5

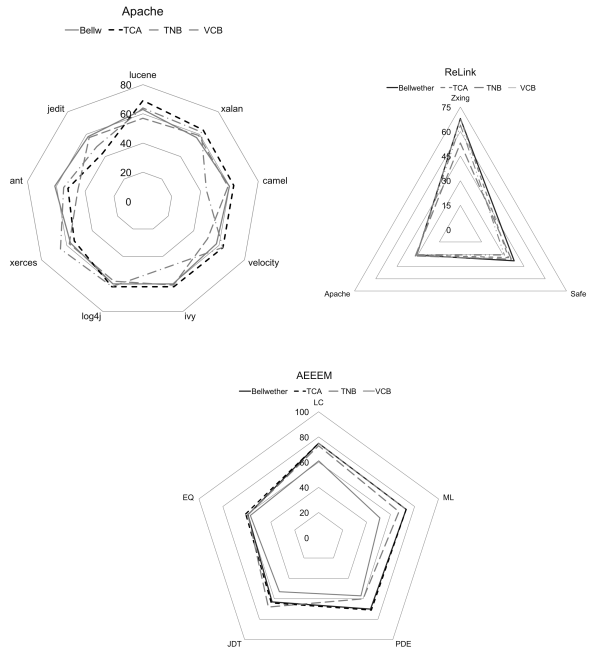

Fig. 10: Defect Datasets: This figure compares the prediction performance of the bellwether dataset (Lucene,Zxing,LC) against other datasets (other rows). *Bellwether Method* against Transfer Learners (columns) for detecting defects. The numerical value seen here are the median G-scores from Equation 2 over 30 repeats where one dataset is used as a source and others are used as targets in a round-robin fashion. Higher values are better and cells highlighted in gray produce the best Scott-Knott ranks. The last row in each community indicate Win/Tie/Loss(W/T/L). *TCA+* is the overall best transfer learner.

2) *Effort Estimation*: When performing effort estimation with functional point analysis, we found that cocomo was the bellwether with remarkably low MMRE scores.
3) *Issue Lifetime:* When predicting for lifetime of issues, we found Node.js to be the bellwether.
4) *Defect datasets:* Finally in the case of defect prediction, Jureczko's bellwether is Lucene; NASA's bellwether is MC; AEEEM's bellwether is LC; and Relink's bellwether is ZXing.

That is, within all the domains studied here, there was always a bellwether dataset for every community. Hence:

Result 1
Our results suggest bellwethers are very prevalent in SE.

RQ2: How well do transfer learners perform across different domains?

We noted in §4 that homogeneous transfer learning was only studied for defect prediction. Also, there was a lack of formal comparison of transfer learner in terms of standardized metrics like Pd, Pf, or G. Hence, here we attempt to perform a comprehensive comparison between these methods applied across the 4 domains under study here (See Figures 8, 9,10, and 11).

In homogeneous transfer learning, the current state-of-the-art is to use TCA+. And indeed our findings from Figure 10 validate this claim. We note that the Bellwether method appears as a close second. However, a similar trend was not noticed the other domains studied here. In estimating code smells for instance, the bellwether method significantly outperforms TCA+. We noticed a similar trend in effort estimation datasets. Here too, Bellwether method outperforms other transfer learners. In issue lifetime estimation, we found (surprisingly) that Transfer Naive Bayes performed much better than other transfer learners.

Our findings indicate that there is no universal "best" transfer learner. But what we do have to report is that a very simple transfer learner like the Bellwether Method can provide competitive results in most cases. This can be used as a benchmark to test other transfer learners. Hence, our answer to this research question is:

Result 2
The performance of transfer learning methods across domains is subject to the constrains of the domain under consideration. No single best transfer learner works across multiple domains.

RQ3: How does the bellwether fare against local models?

Figure 12 compares G-Scores scores of defect predictors built on local models against those built with a bellwether. For this question, we used data from the Apache community since it has the versions required to test older data against newer data.

As seen in the figure, the prediction scores with the bellwether is very encouraging in case of the Apache datasets. In

	Source	**Bellwether**	TCA	TNB
Functional Point Analysis	cocomo	**0.9**	1	2
	coc81	1	5	30
	nasa10	8	20	44
	nasa93	9	62	225
	coc10	196	126	107
	W/T/L	**4/0/1**	0/0/5	1/0/4

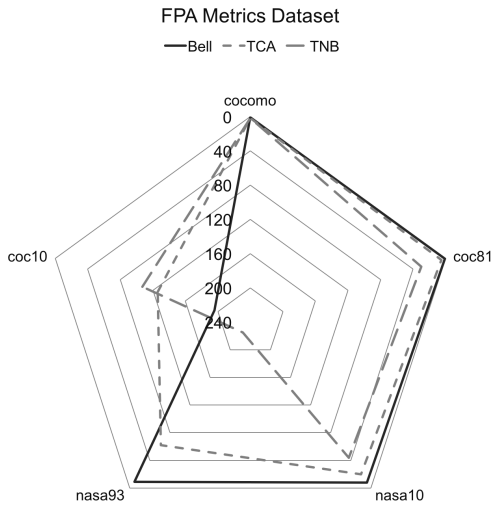

Fig. 11: Effort Estimation: This figure comapres the performance of the bellwether dataset (cocomo) against other datasets (rows) and Transfer Learners (columns) for estimating effort. The numerical value seen are the median MMRE-scores from Equation 3 over 40 repeats. Cells highlighted in gray produce the best Scott-Knott ranks. *Bellwether Method* has the best Win/Tie/Loss ratio.

	Bellw Data & TCA+	Local		
	(Lucene) G-score	Train	Test	G-Score
Ant	70	1.3, 1.4, 1.5, 1.6	1.7	64
Camel	63	1.0, 1.2, 1.4	1.6	57
Xalan	64	2.4, 2.5, 2.6	2.7	55
Xerces	56	1.0, 1.2, 1.3	1.4	38
Ivy	72	1.1, 1.4	2	72
Velocity	59	1.4, 1.5	1.6	59
Log4j	61	1.0, 1.1	1.2	61
Jedit	57	3.2, 4.0, 4.1, 4.2	4.3	62

Fig. 12: Bellwether dataset (Lucene) and TCA+ vs. Local Data. Performance scores are G-scores so *lower* values are *better*. Cells highlighted in gray indicate datasets with superior prediction capability.

all cases, except for Jedit, defect prediction models constructed with the Lucene bellwether performs as well as local data. In some cases like Xerces, Lucene performs much better than local data.

Therefore, the answer to the this research question is:

Result 3
For projects evaluated with the same quality metrics, training a quality prediction model with the Bellwether is just as good as doing so with local data.

RQ4: How much data is required to find the bellwether dataset?

As yet, we do not have a theoretical analysis offering a lower bound for the the number of examples required for finding the bellwether dataset. What we do have is the following empirical observation: all the above results were achieved using SMOTE for sub-sampling (see §5.1); in doing so, a 100 randomly selected interesting modules and 50 randomly selected uninteresting modules we used. Even in cases with as few samples as 25 (code smells dataset) bellwether datasets do exist. Additionally, in cases with limited number of samples and where sub-sampling was not feasible (effort estimation) we discovered bellwethers. Thus:

Result 4
Not much data is required to find bellwether dataset. These bellwethers can be found after projects have discovered as few as 25 samples.

7 Threats to Validity

7.1 Sampling Bias

Sampling bias threatens any classification experiment; what matters in one case may or may not hold in another case. For example, even though we use 100+ open-source datasets in this study which come from several sources, they were all supplied by individuals.

That said, this paper shares this sampling bias problem with every other data mining paper. As researchers, all we can do is document our selection procedure for data (as done in §4) and suggest that other researchers try a broader range of data in future work.

7.2 Learner Bias

For building the quality predictors in this study, we elected to use random forests. We chose this learner because past studies shows that, for prediction tasks, the results were superior to other more complicated algorithms [58] and can act as a baseline for other algorithms.

Apart from this choice, one limitation to our current study is that we have focused here on homogenous transfer learning (where the attributes in source and target are the same). The implications for heterogeneous transfer learning (where the attributes in source an target have different names) are not yet clear. We have some initial results suggesting that a bellwether-like effect occurs when learning across the communities but those results are very preliminary. Hence, for the moment, we would conclude:

- For the homogenous case, we recommend using bellwethers rather than similarity-based transfer learning.
- For the heterogenous case, we recommend using dimensionality tranforms.

7.3 Evaluation Bias

This paper uses one measure of prediction quality, G (see Equation 3). Other quality measures often used in software engineering to quantify the effectiveness of prediction [76] [3] [83] (discussed in §5.2). A comprehensive analysis using these measures may be performed with our replication package.

7.4 Order Bias

With random forest and SMOTE, there is invariably some degree of randomness that is introduced by both the algorithms. Random Forest, as the name suggests, randomly samples the data and constructs trees which it then uses in an ensemble fashion to make predictions.

To mitigate these biases, we run the experiments 30 times (the reruns are equal to 30 in keeping with the central limit theorem). Note that the reported variations over those runs were very small. Hence, we conclude that while order bias is theoretically a problem, it is not a major problem in the particular case of this study.

8 Discussion

The results of this paper cast some doubts on the reusability of the transfer learning results tested on defect prediction. Much of the machinery *does not* translate well to other domains. Bellwethers offer a simpler and a relatively much stable alternative for transfer learning which can motivate development of better transfer learners.

Finally, when reflecting on the findings of this work, there may be three additional questions that arise:

1) *Do bellwethers guarantee permanent conclusion stability?* No- and we should not expect them to. The aim of bellwethers is to *slow*, but do not necessarily *stop*, the pace of new ideas in software engineering (e.g. as in the paper, new quality prediction models). Sometimes, new ideas are essential. Software engineering is a very dynamic field with a high churn in techniques, platforms, developers and tasks. In such a dynamic environment it is important to change with the times. That said, changing *more* than what is necessary is not desirable– hence this paper.
2) *How to detect when bellwether datasets need updating?* The conclusion stability offered by bellwether datasets only lasts as long as the bellwether dataset remains useful. Hence, the bellwether dataset's performance must always be monitoried and, if that performance starts to dip, then seek a new bellwether dataset.
3) *What happens if a set of data has no useful bellwether dataset?* In that case, there are numerous standard transfer learning methods that could be used to import lessons learned from other data [26], [27], [18], [28], [29], [24], [17], [25]. That said, the result here is that all the communities of data explored by this paper had useful bellwether datasets. Hence, we would recommend trying the bellwether method before moving on to more complex methods.

9 Conclusion

In this paper, we have undertaken a detailed study of transfer learners. We have shown that when historical data is limited or not available (e.g. perhaps due the project being in its infancy), developers might seek data from other projects. Our results show that regardless of the sub-domain of software engineering (code smells, effort, defects or issue lifetimes) or granularity of data (see the file, class, file values of Figure 3), there exists a bellwether data set that can be used to train relatively more accurate quality prediction models and this bellwether does not require elaborate data mining methods to discover (just a for-loop around the data sets) and can be found very early in a project's life cycle (after analyzing only a few dozen code modules). Furthermore, we note that previous results on transfer learning do not translate well when applied to domains other than the ones they were designed for. In such cases too, we show that bellwethers are a useful tool to perform accurate and more importantly, stable transfer.

From this, we make the following conclusions. Firstly, this bellwether method is a useful (and very simple) transfer learning method.

- In one of our domains, defect prediction, the initial loop of discovering bellwethers, has to be augmented with TCA+ to perform transfer learning.
- But for the other three domains, we found that merely using the bellwether data without any sort of transfer worked very well.

Hence, from a pragmatic engineering perspective, the primary reason to use bellwethers is the implementation simplicity. In three out of four of our domains, simply wrapping the learner in an evaluation loop was enough to implement a state-of-the-art transfer learner.

Secondly, prior work on homogeneous transfer learning, including some of the authors own papers, may have needless complicated the homogeneous transfer learning process. We strongly recommend that when building increasingly complex automatic methods, researchers should pause and compare their supposedly more sophisticated method against simpler alternatives. Going forward from this paper, we would recommend that the transfer learning community uses bellwethers as a baseline method against which they can test more complex methods.

Acknowledgements

The work is partially funded by NSF awards #1506586 and #1302169.

References

[1] J. Czerwonka, R. Das, N. Nagappan, A. Tarvo, and A. Teterev, "Crane: Failure prediction, change analysis and test prioritization in practice – experiences from windows," in *Software Testing, Verification and Validation (ICST), 2011 IEEE Fourth International Conference on*, march 2011, pp. 357 –366.

[2] T. J. Ostrand, E. J. Weyuker, and R. M. Bell, "Where the bugs are," in *ISSTA '04: Proceedings of the 2004 ACM SIGSOFT international symposium on Software testing and analysis*. New York, NY, USA: ACM, 2004, pp. 86–96.

[3] T. Menzies, A. Dekhtyar, J. Distefano, and J. Greenwald, "Problems with Precision: A Response to "Comments on 'Data Mining Static Code Attributes to Learn Defect Predictors'"," *IEEE Transactions on Software Engineering*, vol. 33, no. 9, pp. 637–640, sep 2007. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4288197

[4] B. Turhan, A. Tosun, and A. Bener, "Empirical evaluation of mixed-project defect prediction models," in *Software Engineering and Advanced Applications (SEAA), 2011 37th EUROMICRO Conference on*. IEEE, 2011, pp. 396–403.

[5] E. Kocaguneli, T. Menzies, A. Bener, and J. Keung, "Exploiting the essential assumptions of analogy-based effort estimation," *IEEE Transactions on Software Engineering*, vol. 28, pp. 425–438, 2012, available from http://menzies.us/pdf/11teak.pdf.

[6] T. Savor, M. Douglas, M. Gentili, L. Williams, K. Beck, and M. Stumm, "Continuous deployment at facebook and oanda," in *Proceedings of the 38th International Conference on Software Engineering Companion*. ACM, 2016, pp. 21–30.

[7] M. Linares-Vásquez, G. Bavota, C. Bernal-Cárdenas, R. Oliveto, M. Di Penta, and D. Poshyvanyk, "Mining energy-greedy api usage patterns in android apps: an empirical study," in *Proceedings of the 11th Working Conference on Mining Software Repositories*. ACM, 2014, pp. 2–11.

[8] A. Begel and T. Zimmermann, "Analyze this! 145 questions for data scientists in software engineering," in *Proceedings of the 36th International Conference on Software Engineering (ICSE 2014)*. ACM, June 2014. [Online]. Available: http://research.microsoft.com/apps/pubs/default.aspx?id=208800

[9] C. Theisen, K. Herzig, P. Morrison, B. Murphy, and L. Williams, "Approximating attack surfaces with stack traces," in *ICSE'15*, 2015.

[10] T. Zimmermann and T. Menzies, "Software analytics: So what?" *IEEE Software*, vol. 30, no. 4, pp. 0031–37, 2013.

[11] C. Bird, T. Menzies, and T. Zimmermann, *The Art and Science of Analyzing Software Data*. Elsevier, 2015.

[12] F. Sarro, A. Petrozziello, and M. Harman, "Multi-objective software effort estimation," in *38th International Conference on Software Engineering (ICSE'16)*, no. to appear. ACM, 2016.

[13] M. Choetkiertikul, H. K. Dam, T. Tran, and A. Ghose, "Predicting delays in software projects using networked classification (t)," in *Automated Software Engineering (ASE), 2015 30th IEEE/ACM International Conference on*. IEEE, 2015, pp. 353–364.

[14] F. Rahman, S. Khatri, E. T. Barr, and P. Devanbu, "Comparing static bug finders and statistical prediction," in *Proceedings of the 36th International Conference on Software Engineering*. ACM, 2014, pp. 424–434.

[15] Y. Shin and L. Williams, "Can traditional fault prediction models be used for vulnerability prediction?" *Empirical Software Engineering*, vol. 18, no. 1, pp. 25–59, 2013.

[16] B. Turhan, T. Menzies, A. Bener, and J. Distefano, "On the relative value of cross-company and within-company data for defect prediction," *Empirical Software Engineering*, vol. 68, no. 2, pp. 278–290, 2009, available from http://menzies.us/pdf/08ccwc.pdf.

[17] J. Nam and S. Kim, "Heterogeneous defect prediction," in *Proc. 2015 10th Jt. Meet. Found. Softw. Eng. - ESEC/FSE 2015*. New York, New York, USA: ACM Press, 2015, pp. 508–519. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2786805.2786814

[18] Z. He, F. Peters, T. Menzies, and Y. Yang, "Learning from open-source projects: An empirical study on defect prediction," in *Empirical Software Engineering and Measurement, 2013 ACM/IEEE International Symposium on*. IEEE, 2013, pp. 45–54.

[19] A. Tosun, A. Bener, and R. Kale, "AI-based software defect predictors: Applications and benefits in a case study," in *Twenty-Second IAAI Conference on Artificial Intelligence*, 2010.

[20] A. Tosun, A. Bener, and B. Turhan, "Practical considerations of deploying ai in defect prediction: A case study within the Turkish telecommunication industry," in *PROMISE'09*, 2009.

[21] M. Lumpe, R. Vasa, T. Menzies, R. Rush, and R. Turhan, "Learning better inspection optimization policies," *International Journal of Software Engineering and Knowledge Engineering*, vol. 21, no. 45, pp. 725–753, 2011.

[22] C. Bird, A. Bachmann, E. Aune, J. Duffy, A. Bernstein, V. Filkov, and P. Devanbu, "Fair and balanced?: bias in bug-fix datasets," in *Proceedings of the the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering*. ACM, 2009, pp. 121–130.

[23] F. Rahman and P. Devanbu, "How, and why, process metrics are better," in *2013 35th International Conference on Software Engineering (ICSE)*, May 2013, pp. 432–441.

[24] J. Nam, S. J. Pan, and S. Kim, "Transfer defect learning," in *Proceedings - International Conference on Software Engineering*, 2013, pp. 382–391.

[25] X. Jing, F. Wu, X. Dong, F. Qi, and B. Xu, "Heterogeneous Cross-Company Defect Prediction by Unified Metric Representation and CCA-Based Transfer Learning Categories and Subject Descriptors," *Proceeding of the 10th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE 2015)*, pp. 496–507, 2015.

[26] E. Kocaguneli and T. Menzies, "How to find relevant data for effort estimation?" in *Empirical Software Engineering and Measurement (ESEM), 2011 International Symposium on*. IEEE, 2011, pp. 255–264.

[27] E. Kocaguneli, T. Menzies, and E. Mendes, "Transfer learning in effort estimation," *Empirical Software Engineering*, vol. 20, no. 3, pp. 813–843, jun 2015. [Online]. Available: http://link.springer.com/10.1007/s10664-014-9300-5

[28] B. Turhan, T. Menzies, A. B. Bener, and J. Di Stefano, "On the relative value of cross-company and within-company data for defect prediction," *Empirical Software Engineering*, vol. 14, no. 5, pp. 540–578, 2009.

[29] F. Peters, T. Menzies, and L. Layman, "LACE2: Better privacy-preserving data sharing for cross project defect prediction," in *Proceedings - International Conference on Software Engineering*, vol. 1, 2015, pp. 801–811.

[30] F. Rahman, D. Posnett, and P. Devanbu, "Recalling the "imprecision" of cross-project defect prediction," in *Proceedings of the ACM SIGSOFT 20th International Symposium on the Foundations of Software Engineering*, ser. FSE '12. New York, NY, USA: ACM, 2012, pp. 61:1–61:11. [Online]. Available: http://doi.acm.org/10.1145/2393596.2393669

[31] R. Krishna, T. Menzies, and W. Fu, "Too Much Automation? The Bellwether Effect and Its Implications for Transfer Learning," in *ASE'16*, 2016.

[32] Y. Ma, G. Luo, X. Zeng, and A. Chen, "Transfer learning for cross-company software defect prediction," *Information and Software Technology*, vol. 54, no. 3, pp. 248–256, 2012.

[33] D. Ryu, O. Choi, and J. Baik, "Value-cognitive boosting with a support vector machine for cross-project defect prediction," *Empir. Softw. Eng.*, vol. 21, no. 1, pp. 43–71, feb 2016. [Online]. Available: http://link.springer.com/10.1007/s10664-014-9346-4

[34] T. Zimmermann, N. Nagappan, H. Gall, E. Giger, and B. Murphy, "Cross-project defect prediction: a large scale experiment on data vs. domain vs. process," in *Proceedings of the the 7th joint meeting of the European software engineering conference and the ACM SIGSOFT symposium on The foundations of software engineering*. ACM, 2009, pp. 91–100.

[35] F. Zhang, A. Mockus, I. Keivanloo, and Y. Zou, "Towards building a universal defect prediction model with rank transformed predictors," *Empirical Software Engineering*, pp. 1–39, 2015. [Online]. Available: http://dx.doi.org/10.1007/s10664-015-9396-2

[36] T. Hall, S. Beecham, D. Bowes, D. Gray, and S. Counsell, "A systematic literature review on fault prediction performance in software engineering," *IEEE Transactions on Software Engineering*, vol. 38, no. 6, pp. 1276–1304, Nov 2012.

[37] R. Krishna, T. Menzies, and L. Layman, "Less is More: Minimizing Code Reorganization using XTREE," *CoRR*, vol. abs/1609.03614, 2016. [Online]. Available: http://arxiv.org/abs/1609.03614

[38] M. Fowler, K. Beck, J. Brant, W. Opdyke, and D. Roberts, *Refactoring: Improving the Design of Existing Code*. Boston, MA, USA: Addison-Wesley Longman, 1999.

[39] A. Yamashita and S. Counsell, "Code smells as system-level indicators of maintainability: An empirical study," *Journal of Systems and Software*, vol. 86, no. 10, pp. 2639–2653, 2013.

[40] A. Yamashita and L. Moonen, "Exploring the impact of inter-smell relations on software maintainability: An empirical study," in *Proceedings of the 2013 International Conference on Software Engineering*. IEEE Press, 2013, pp. 682–691.

[41] N. Zazworka, M. A. Shaw, F. Shull, and C. Seaman, "Investigating the impact of design debt on software quality," in *Proceedings of the 2nd Workshop on Managing Technical Debt*. ACM, 2011, pp. 17–23.

[42] F. A. Fontana, P. Braione, and M. Zanoni, "Automatic detection of bad smells in code: An experimental assessment." *Journal of Object Technology*, vol. 11, no. 2, pp. 5–1, 2012.

[43] J. Kreimer, "Adaptive Detection of Design Flaws," *Electronic Notes in Theoretical Computer Science*, vol. 141, no. 4, pp. 117–136, 2005. [Online]. Available: //www.sciencedirect.com/science/article/pii/S1571066105051844

[44] F. Khomh, S. Vaucher, Y. G. Guhneuc, and H. Sahraoui, "A bayesian approach for the detection of code and design smells," in *2009 Ninth International Conference on Quality Software*, Aug 2009, pp. 305–314.

[45] F. Khomh, S. Vaucher, Y.-G. Guhneuc, and H. Sahraoui, "Bdtex: A gqm-based bayesian approach for the detection of antipatterns," *Journal of Systems and Software*, vol. 84, no. 4, pp. 559 – 572, 2011, the Ninth International Conference on Quality Software. [Online]. Available: //www.sciencedirect.com/science/article/pii/S0164121210003225

[46] J. Yang, K. Hotta, Y. Higo, H. Igaki, and S. Kusumoto, "Filtering clones for individual user based on machine learning analysis," in *2012 6th International Workshop on Software Clones (IWSC)*, June 2012, pp. 76–77.

[47] F. Arcelli Fontana, M. V. Mäntylä, M. Zanoni, and A. Marino, "Comparing and experimenting machine learning techniques for code smell detection," *Empir. Softw. Eng.*, vol. 21, no. 3, pp. 1143–1191, jun 2016. [Online]. Available: http://dx.doi.org/10.1007/s10664-015-9378-4http://link.springer.com/10.1007/s10664-015-9378-4

[48] E. Tempero, C. Anslow, J. Dietrich, T. Han, J. Li, M. Lumpe, H. Melton, and J. Noble, "Qualitas corpus: A curated collection of java code for empirical studies," in *2010 Asia Pacific Software Engineering Conference (APSEC2010)*, Dec. 2010, pp. 336–345.

[49] L. D. Panjer, "Predicting eclipse bug lifetimes," in *Proceedings of the Fourth International Workshop on Mining Software Repositories*, ser. MSR '07. Washington, DC, USA: IEEE Computer Society, 2007, pp. 29–. [Online]. Available: http://dx.doi.org/10.1109/MSR.2007.25

[50] E. Giger, M. Pinzger, and H. Gall, "Predicting the fix time of bugs," in *Proceedings of the 2Nd International Workshop on Recommendation Systems for Software Engineering*, ser. RSSE '10. New York, NY, USA: ACM, 2010, pp. 52–56. [Online]. Available: http://doi.acm.org/10.1145/1808920.1808933

[51] H. Zhang, L. Gong, and S. Versteeg, "Predicting bug-fixing time: An empirical study of commercial software projects," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 1042–1051. [Online]. Available: http://dl.acm.org/citation.cfm?id=2486788.2486931

[52] M. Rees-jones, M. Martin, C. College, and T. Menzies, "Better Predictors for Issue Lifetime," , pp. 1–8, 2017.

[53] R. Kikas, M. Dumas, and D. Pfahl, "Using dynamic and contextual features to predict issue lifetime in github projects," in *Proceedings of the 13th International Conference on Mining Software Repositories*, ser. MSR '16. New York, NY, USA: ACM, 2016, pp. 291–302. [Online]. Available: http://doi.acm.org/10.1145/2901739.2901751

[54] B. A. Kitchenham, E. Mendes, and G. H. Travassos, "Cross versus within-company cost estimation studies: A systematic review," *IEEE Trans. Softw. Eng.*, vol. 33, no. 5, pp. 316–329, May 2007. [Online]. Available: http://dx.doi.org/10.1109/TSE.2007.1001

[55] Y. Yang, L. Xie, Z. He, Q. Li, V. Nguyen, B. Boehm, and R. Valerdi, "Local bias and its impacts on the performance of parametric estimation models," in *Proceedings of the 7th International Conference on Predictive Models in Software Engineering - Promise '11*. New York, New York, USA: ACM Press, 2011, pp. 1–10. [Online]. Available: http://dl.acm.org/citation.cfm?doid=2020390.2020404

[56] T. Menzies, Y. Yang, G. Mathew, B. Boehm, and J. Hihn, "Negative results for software effort estimation," *Empirical Software Engineering*, pp. 1–26, 2016. [Online]. Available: http://dx.doi.org/10.1007/s10664-016-9472-2

[57] B. W. Boehm *et al.*, *Software engineering economics*. Prentice-hall Englewood Cliffs (NJ), 1981, vol. 197.

[58] S. Lessmann, B. Baesens, C. Mues, and S. Pietsch, "Benchmarking Classification Models for Software Defect Prediction: A Proposed Framework and Novel Findings," *IEEE Trans. Softw. Eng.*, vol. 34, no. 4, pp. 485–496, jul 2008. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4527256

[59] M. Lowry, M. Boyd, and D. Kulkami, "Towards a theory for integration of mathematical verification and empirical testing," in *Automated Software Engineering, 1998. Proceedings. 13th IEEE International Conference on*. IEEE, 1998, pp. 322–331.

[60] N. Nagappan and T. Ball, "Static analysis tools as early indicators of pre-release defect density," in *ICSE 2005, St. Louis*, 2005.

[61] T. Menzies, D. Raffo, S. on Setamanit, Y. Hu, and S. Tootoonian, "Model-based tests of truisms," in *Proceedings of IEEE ASE 2002*, 2002, available from http://menzies.us/pdf/02truisms.pdf.

[62] C. Lewis, Z. Lin, C. Sadowski, X. Zhu, R. Ou, and E. J. Whitehead Jr., "Does bug prediction support human developers? findings from a google case study," in *Proceedings of the 2013 International Conference on Software Engineering*, ser. ICSE '13. Piscataway, NJ, USA: IEEE Press, 2013, pp. 372–381. [Online]. Available: http://dl.acm.org/citation.cfm?id=2486788.2486838

[63] S. Rakitin, *Software Verification and Validation for Practitioners and Managers, Second Edition*. Artech House, 2001.

[64] T. Menzies, J. Greenwald, and A. Frank, "Data mining static code attributes to learn defect predictors," *IEEE Transactions on Software Engineering*, January 2007, available from http://menzies.us/pdf/06learnPredict.pdf.

[65] F. Shull, V. B. ad B. Boehm, A. Brown, P. Costa, M. Lindvall, D. Port, I. Rus, R. Tesoriero, and M. Zelkowitz, "What we have learned about fighting defects," in *Proceedings of 8th International Software Metrics Symposium, Ottawa, Canada*, 2002, pp. 249–258.

[66] M. Fagan, "Design and code inspections to reduce errors in program development," *IBM Systems Journal*, vol. 15, no. 3, 1976.

[67] M. D'Ambros, M. Lanza, and R. Robbes, "Evaluating defect prediction approaches: a benchmark and an extensive comparison," *Empir. Softw. Eng.*, vol. 17, no. 4-5, pp. 531–577, aug 2012. [Online]. Available: http://link.springer.com/10.1007/s10664-011-9173-9

[68] R. Wu, H. Zhang, S. Kim, and S.-C. Cheung, "ReLink," in *Proc. 19th ACM SIGSOFT Symp. 13th Eur. Conf. Found. Softw. Eng. - SIGSOFT/FSE '11*. New York, New York, USA: ACM Press, 2011, p. 15.

[69] V. R. Basili, L. C. Briand, and W. L. Melo, "A validation of object-oriented design metrics as quality indicators," *Software Engineering, IEEE Transactions on*, vol. 22, no. 10, pp. 751–761, 1996.

[70] N. Ohlsson and H. Alberg, "Predicting fault-prone software modules in telephone switches," *Software Engineering, IEEE Transactions on*, vol. 22, no. 12, pp. 886–894, 1996.

[71] S. Kim, H. Zhang, R. Wu, and L. Gong, "Dealing with noise in defect prediction," in *Software Engineering (ICSE), 2011 33rd International Conference on*. IEEE, 2011, pp. 481–490.

[72] M. Jureczko and L. Madeyski, "Towards identifying software project clusters with regard to defect prediction," in *Proc. 6th Int. Conf. Predict. Model. Softw. Eng. - PROMISE '10*. New York, New York, USA: ACM Press, 2010, p. 1. [Online]. Available: http://portal.acm.org/citation.cfm?doid=1868328.1868342

[73] L. Breiman, "Random forests," *Machine learning*, pp. 5–32, 2001. [Online]. Available: http://link.springer.com/article/10.1023/A:1010933404324

[74] L. Pelayo and S. Dick, "Applying Novel Resampling Strategies To Software Defect Prediction," in *NAFIPS 2007 - 2007 Annu. Meet. North Am. Fuzzy Inf. Process. Soc.* IEEE, jun 2007, pp. 69–72. [Online]. Available: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=4271036

[75] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "SMOTE: Synthetic minority over-sampling technique," *J. Artif. Intell. Res.*, vol. 16, 2002.

[76] Y. Ma and B. Cukic, "Adequate and precise evaluation of quality models in software engineering studies," in *Predictor Models in Software Engineering, 2007. PROMISE'07: ICSE Workshops 2007. International Workshop on*, May 2007, pp. 1–1.

[77] N. Mittas and L. Angelis, "Ranking and clustering software cost estimation models through a multiple comparisons algorithm," *IEEE Trans. Software Eng.*, vol. 39, no. 4, pp. 537–551, 2013.

[78] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*, ser. Mono. Stat. Appl. Probab. London: Chapman and Hall, 1993.

[79] M. J. Shepperd and S. G. MacDonell, "Evaluating prediction systems in software project estimation," *Information & Software Technology*, vol. 54, no. 8, pp. 820–827, 2012.

[80] V. B. Kampenes, T. Dybå, J. E. Hannay, and D. I. K. Sjøberg, "A systematic review of effect size in software engineering experiments," *Information & Software Technology*, vol. 49, no. 11-12, pp. 1073–1086, 2007.

[81] E. Kocaguneli, T. Zimmermann, C. Bird, N. Nagappan, and T. Menzies, "Distributed development considered harmful?" in *Proceedings - International Conference on Software Engineering*, 2013, pp. 882–890.

[82] A. Arcuri and L. Briand, "A practical guide for using statistical tests to assess randomized algorithms in software engineering," in *ICSE'11*, 2011, pp. 1–10.

[83] W. Fu, T. Menzies, and X. Shen, "Tuning for software analytics: Is it really necessary?" *Information and Software Technology*, vol. 76, pp. 135 – 146, 2016. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0950584916300738